

49

Notice of Allowability	Application No.	Applicant(s)	
	10/721,596	FAIR, ROBERT L.	
	Examiner	Art Unit	
	Navneet K. Ahluwalia	2166	

-- The MAILING DATE of this communication appears on the cover sheet with the correspondence address--

All claims being allowable, PROSECUTION ON THE MERITS IS (OR REMAINS) CLOSED in this application. If not included herewith (or previously mailed), a Notice of Allowance (PTOL-85) or other appropriate communication will be mailed in due course. **THIS NOTICE OF ALLOWABILITY IS NOT A GRANT OF PATENT RIGHTS.** This application is subject to withdrawal from issue at the initiative of the Office or upon petition by the applicant. See 37 CFR 1.313 and MPEP 1308.

1. ☒ This communication is responsive to 09/12/2007.
2. ☒ The allowed claim(s) is/are 1 - 3, 5 - 14, 16 - 18, 29 - 43 and 45 - 47 (renumbered as 1 - 34).
3. ☐ Acknowledgment is made of a claim for foreign priority under 35 U.S.C. § 119(a)-(d) or (f).
 - a) ☐ All b) ☐ Some* c) ☐ None of the:
 1. ☐ Certified copies of the priority documents have been received.
 2. ☐ Certified copies of the priority documents have been received in Application No. _____.
 3. ☐ Copies of the certified copies of the priority documents have been received in this national stage application from the International Bureau (PCT Rule 17.2(a)).
 - * Certified copies not received: _____.

Applicant has THREE MONTHS FROM THE "MAILING DATE" of this communication to file a reply complying with the requirements noted below. Failure to timely comply will result in ABANDONMENT of this application.
THIS THREE-MONTH PERIOD IS NOT EXTENDABLE.

4. ☐ A SUBSTITUTE OATH OR DECLARATION must be submitted. Note the attached EXAMINER'S AMENDMENT or NOTICE OF INFORMAL PATENT APPLICATION (PTO-152) which gives reason(s) why the oath or declaration is deficient.
5. ☐ CORRECTED DRAWINGS (as "replacement sheets") must be submitted.
 - (a) ☐ including changes required by the Notice of Draftsperson's Patent Drawing Review (PTO-948) attached
 - 1) ☐ hereto or 2) ☐ to Paper No./Mail Date _____.
 - (b) ☐ including changes required by the attached Examiner's Amendment / Comment or in the Office action of Paper No./Mail Date _____.

Identifying indicia such as the application number (see 37 CFR 1.84(c)) should be written on the drawings in the front (not the back) of each sheet. Replacement sheet(s) should be labeled as such in the header according to 37 CFR 1.121(d).
6. ☐ DEPOSIT OF and/or INFORMATION about the deposit of BIOLOGICAL MATERIAL must be submitted. Note the attached Examiner's comment regarding REQUIREMENT FOR THE DEPOSIT OF BIOLOGICAL MATERIAL.

Attachment(s)

- | | |
|--|--|
| 1. <input type="checkbox"/> Notice of References Cited (PTO-892) | 5. <input type="checkbox"/> Notice of Informal Patent Application |
| 2. <input type="checkbox"/> Notice of Draftsperson's Patent Drawing Review (PTO-948) | 6. <input checked="" type="checkbox"/> Interview Summary (PTO-413),
Paper No./Mail Date <u>20070914</u> . |
| 3. <input type="checkbox"/> Information Disclosure Statements (PTO/SB/08),
Paper No./Mail Date _____ | 7. <input checked="" type="checkbox"/> Examiner's Amendment/Comment |
| 4. <input type="checkbox"/> Examiner's Comment Regarding Requirement for Deposit
of Biological Material | 8. <input type="checkbox"/> Examiner's Statement of Reasons for Allowance |
| | 9. <input type="checkbox"/> Other _____. |

W. K. Ahluwalia
 Mohammad Ali
 SPE 2169

DETAILED ACTION

1. This communication is in response to the applicant-initiated interview on 09/12/2007.

After a search and thorough examination of the present application and in light of the prior art made of records, claims 1 – 3, 5 – 14, 16 – 18, 29 – 43 and 45 – 47 (renumbered as 1 – 34) are allowed.

EXAMINER'S AMENDMENT

2. An examiner's amendment to the record appears below. Should the changes and/or additions be unacceptable to applicant, an amendment may be filed as provided by 37 CFR 1.312. To ensure consideration of such an amendment, it **MUST** be submitted no later than the payment of the issue fee.

Authorization for this examiner's amendment was given in a telephone interview with Shannen C. Delaney, Reg. No. 51,695 on 09/12/2007.

The application has been amended as follows:

The claims read as:

1. (Currently Amended) A method for a storage operating system implemented in a storage system to concurrently perform readahead operations for a plurality of different read streams established in one or more files, directories, vdisks or luns stored in the storage system, the method comprising:

allocating at least one readset data structure ("readset") for each of the one or more files, directories, vdisks or luns in which the plurality of different read streams is established, wherein the number of readsets allocated for each file, directory, vdisk or lun depends on the size of that file, directory, vdisk or lun;

receiving a client read request at the storage system, the client read request indicating client-requested data for the storage operating system to retrieve from a file, directory, vdisk or lun stored in the storage system;

determining whether the received client read request matches any of the plurality of readsets allocated for the file, directory, vdisk or lun containing the client-requested data; performing readahead operations in accordance with a set of readahead metadata stored in an associated readset that is determined to match the received client read request, wherein the readahead metadata describes the associated readset; and

if the received client read request does not match any of the readsets allocated for the file, directory, vdisk or lun containing the client-requested data, then performing the steps:

identifying the received client read request as being the first read request in a new read stream;

generating a set of readahead metadata associated with the new read stream;

selecting for reuse one of the readsets allocated for the file, directory, vdisk or lun containing the client-requested data; and

storing the generated set of readahead metadata associated with the new read stream in the readset selected for reuse.

2. (Previously Presented) The method of claim 1, further comprising:

generating a separate set of readahead metadata for each of the plurality of different read streams; and

storing each generated set of readahead metadata in a different readset allocated for the file, directory, vdisk or lun in which the read stream associated with the generated set of readahead metadata is established.

3. (Original) The method of claim 1, further comprising:

initializing each allocated readset to store a predetermined set of values.

4. (Cancelled)
5. (Previously Presented) The method of claim 2, wherein the number of readsets allocated for a file, directory, vdisk or lun is dynamically increased as the size of that file, directory, vdisk or lun is increased.
6. (Original) The method of claim 1, wherein a first readset is determined to match the received client read request if the first readset stores a set of readahead metadata associated with a read stream that is extended by the client-requested data.
7. (Original) The method of claim 1, wherein a second readset is determined to match the received client read request when the client-requested data is located within a predetermined fuzzy range associated with the second readset.
8. (Original) The method of claim 7, wherein the fuzzy range is derived based on a multiple of a number of client-requested data blocks specified in the received client read request.
9. (Original) The method of claim 7, wherein the fuzzy range extends in both a forward direction and a backward direction in relation to a last data block retrieved in a read stream associated with the second readset.
10. (Original) The method of claim 1, wherein a third readset is determined to match the received client read request if the third readset is determined to be unused.
11. (Original) The method of claim 10, wherein the third readset is determined to be unused when a level value stored in the third readset equals a special indicator value.

Art Unit: 2166

12. (Original) The method of claim 1, wherein readahead operations are not performed if the storage operating system determines that the file, directory, vdisk or lun containing the client-requested data is accessed using a random access style.

13. (Original) The method of claim 12, wherein a DAFS cache hint included in the received client read request indicates that the file, directory, vdisk or lun containing the client-requested data is accessed using a random access style.

14. (Original) The method of claim 1, wherein readahead operations are not performed unless:
(i) a readset is determined to match the received client read request; and
(ii) the matching readset stores a set of readahead metadata associated with a read stream that is extended by the client-requested data past a predetermined data block or memory address.

15. (Cancelled)

16. (Currently Amended) The method of claim 1, wherein the readset selected for reuse stores a level value that is less than or equal to level values stored in each of the other readsets associated with the file, directory, vdisk or lun containing the client-requested data.

17. (Original) The method of claim 1, wherein the client read request received at the storage system is a file-based client read request.

18. (Original) The method of claim 1, wherein the client read request received at the storage system is a block-based client read request.

19-28 (Cancelled)

29. (Currently Amended) A storage system that employs a storage operating system to concurrently perform readahead operations for a plurality of different read streams established in

one or more files, directories, vdisks or luns stored in the storage system, the storage system comprising:

means for allocating at least one readset data structure ("readset") for each of the one or more files, directories, vdisks or luns in which the plurality of different read streams is established, wherein the number of readsets allocated for each file, directory, vdisk or lun depends on the size of that file, directory, vdisk or lun;

means for receiving a client read request at the storage system, the client read request indicating client-requested data for the storage operating system to retrieve from a file, directory, vdisk or lun stored in the storage system;

means for determining whether the received client read request matches any of the plurality of readsets" allocated for the file, directory, vdisk or lun containing the client-requested data;

means for performing readahead operations in accordance with a set of readahead metadata stored in an associated readset that is determined to match the received client read request, wherein the readahead metadata describes the associated readset; and

if the received client read request does not match any of the readsets allocated for the file, directory, vdisk or lun containing the client-requested data, then means for performing:

means for identifying the received client read request as being the first read request in a new read stream;

means for generating a set of readahead metadata associated with the new read stream;

means for selecting for reuse one of the readsets allocated for the file, directory, vdisk or lun containing the client-requested data; and

means for storing the generated set of readahead metadata associated with the new read stream in the readset selected for reuse.

30. (Currently Amended) A computer-readable media comprising instructions for execution in a processor for the practice of a method for a storage operating system implemented in a storage system to concurrently perform readahead operations for a plurality of different read streams established in one or more files, directories, vdisks or luns stored in the storage system, the method comprising:

allocating at least one readset data structure ("readset") for each of the one or more files, directories, vdisks or luns in which the plurality of different read streams is established, wherein the number of readsets allocated for each file, directory, vdisk or lun depends on the size of that file, directory, vdisk or lun;

receiving a client read request at the storage system, the client read request indicating client-requested data for the storage operating system to retrieve from a file, directory, vdisk or lun stored in the storage system;

determining whether the received client read request matches any of the plurality of readsets allocated for the file, directory, vdisk or lun containing the client-requested data;

performing readahead operations in accordance with a set of readahead metadata stored in an associated readset that is determined to match the received client read request, wherein the readahead metadata describes the associated readset; and

if the received client read request does not match any of the readsets allocated for the file, directory, vdisk or lun containing the client-requested data, then performing the steps:

identifying the received client read request as being the first read request in a new read stream;

generating a set of readahead metadata associated with the new read stream;

selecting for reuse one of the readsets allocated for the file, directory, vdisk or lun containing the client-requested data; and

storing the generated set of readahead metadata associated with the new read stream in the readset selected for reuse.

31. (Currently Amended) A method for a storage operating system implemented in a storage system to concurrently perform readahead operations for a plurality of different read streams established in one or more files stored in the storage system, comprising:

allocating at least one read set data structure ("readset") for each of the one or more files,

directories, vdisks or luns in which the plurality of different read streams is established wherein the number of readsets allocated for each file depends on the size of that file;

generating a separate set of readahead metadata for each of the plurality of different read streams; and

storing each generated set of readahead metadata in a different readset allocated for the file in which the read stream associated with the generated set of readahead metadata is established; receiving a client read request at the storage system, the client read request indicating client-requested data for the storage operating system to retrieve from a file, stored in the storage system;

determining whether the received client read request matches any of a plurality of readsets allocated for the file containing the client-requested data; and performing readahead operations in accordance with a set of readahead metadata stored in a readset that is determined to match the received client read request; and if the received client read request does not match any of the readsets allocated for the file, directory, vdisk or lun containing the client-requested data, then performing the steps: identifying the received client read request as being the first read request in a new read stream; generating a set of readahead metadata associated with the new read stream; selecting for reuse one of the readsets allocated for the file, directory, vdisk or lun containing the client-requested data; and storing the generated set of readahead metadata associated with the new read stream in the readset selected for reuse.

32. (Previously Presented) The method of claim 31, wherein the file is broad term describing either a file, directory, vdisk or lun.

33. (Previously Presented) The method of claim 31, further comprising: initializing each allocated readset to store a predetermined set of values.

34. (Previously Presented) The method of claim 31, wherein the number of readsets allocated for a file is dynamically increased as the size of that file is increased.

35. (Previously Presented) The method of claim 31, wherein a first readset is determined to match the received client read request if the first readset stores a set of readahead metadata associated with a read stream that is extended by the client-requested data.

36. (Previously Presented) The method of claim 31, wherein a second readset is determined to match the received client read request when the client-requested data is located within a predetermined fuzzy range associated with the second readset.

37. (Previously Presented) The method of claim 36, wherein the fuzzy range is derived based on a multiple of a number of client-requested data blocks specified in the received client read request.

38. (Previously Presented) The method of claim 36, wherein the fuzzy range extends in both a forward direction and a backward direction in relation to a last data block retrieved in a read stream associated with the second readset.

39. (Previously Presented) The method of claim 31, wherein a third readset is determined to match the received client read request if the third readset is determined to be unused.

40. (Previously Presented) The method of claim 39, wherein the third readset is determined to be unused when a level value stored in the third readset equals a special indicator value.

41. (Previously Presented) The method of claim 31, wherein readahead operations are not performed if the storage operating system determines that the file, directory, vdisk or lun containing the client-requested data is accessed using a random access style.

42. (Previously Presented) The method of claim 41, wherein a DAFS cache hint included in the received client read request indicates that the file, directory, vdisk or lun containing the client-requested data is accessed using a random access style.

43. (Previously Presented) The method of claim 31, wherein readahead operations are not performed unless:

(i) a readset is determined to match the received client read request; and

Art Unit: 2166

(ii) the matching readset stores a set of readahead metadata associated with a read stream that is extended by the client-requested data past a predetermined data block or memory address.

44. (Cancelled)

45. (Currently Amended) The method of claim 31, wherein the readset selected for reuse stores a level value that is less than or equal to level values stored in each of the other readsets associated with the file, directory, vdisk or lun containing the client-requested data.

46. (Previously Presented) The method of claim 31, wherein the client read request received at the storage system is a file-based client read request.

47. (Previously Presented) The method of claim 31, wherein the client read request received at the storage system is a block-based client read request.

Contact Information

Any inquiry concerning this communication or earlier communications from the examiner should be directed to Navneet K. Ahluwalia whose telephone number is 571-272-5636.

If attempts to reach the examiner by telephone are unsuccessful, the examiner's supervisor, Alam T. Hosain can be reached on 571-272-3978. The fax phone number for the organization where this application or proceeding is assigned is 571-273-8300.

Information regarding the status of an application may be obtained from the Patent Application Information Retrieval (PAIR) system. Status information for published applications may be obtained from either Private PAIR or Public PAIR. Status information for unpublished applications is available through Private PAIR only. For more information about the PAIR system, see <http://pair-direct.uspto.gov>. Should you have questions on access to the Private PAIR system, contact the Electronic Business Center (EBC) at 866-217-9197 (toll-free). If you would like assistance from a USPTO Customer Service Representative or access to the automated information system, call 800-786-9199 (IN USA OR CANADA) or 571-272-1000.

Navneet

Navneet K. Ahluwalia
Examiner
Art Unit 2166

Dated: 09/14/2007

urh al
SPE 2169